



FusionGrid: Bringing the Sun to the Earth

A MICS/SciDAC National Collaboratory Pilot Project




Presented by
David P. Schissel
Lead-PI

at
SC2004
November 6-12, 2004
Pittsburgh, PA



ACKNOWLEDGMENT


- Work is supported by the Department of Energy
 - A MICS/SciDAC National Collaboratory Project
 - Office of Fusion Energy Sciences
- The National Fusion Collaboratory Team
 - ANL Distributed Systems Lab – I. Foster, K. Keahey
 - ANL Futures Lab – T. Legget, M. Papka, R. Stevens
 - General Atomics – G. Abla, J. Burruss, D. Schissel
 - LBL Distributed Systems – M. Thompson
 - MIT PSFC – M. Fredian, M. Greenwald, J. Stillerman
 - Princeton Computer Science – K. Li, G. Wallace
 - Princeton Plasma Physics – E. Feibush, D. McCune, L. Randerson
 - Scientific Computing & Imaging Institute – C. Johnson, A. Sanderson



PRESENTATION'S KEY POINTS

- Collaborative technology critical to the success of the FE program

- Experimental: Fewer, larger machines in future (KSTAR, ITER)
- Includes design, engineering and construction phase
- Computation: Toward an integrated simulation (FSP)



- The collaborative control room is being realized

- Secure computational resources scheduled as required
- Rapidly compare experimental data to simulation results
- Share individual results with the group via shared displays
- Fully engaged remote scientists with audio, video, shared displays


- Work is still needed to make for routine use

- Ease-of-use, but not neglecting security (“security with transparency”)
- Interoperability of technologies

FUSION IS THE PROCESS OF THE STARS AND OUR SUN

- Joining two light nuclei ($E=mc^2$)

- Pickup truck fusion fuel = 21,000 railroad cars of coal



- Like charges repel

- High velocity (temperature)

- High temp removes electrons

- Plasma

- Charge particles follow magnetic field lines

- Magnetic donut = tokamak

- For now, bigger is better


- Knowledge will lead to smaller devices



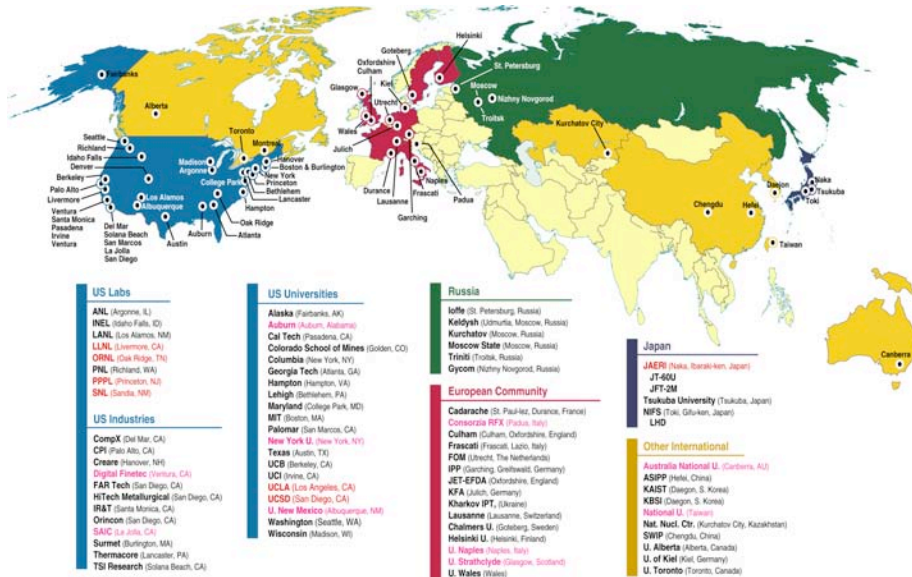
Long term goal to develop a reliable energy system that is environmentally and economically sustainable

FUSION: NON-GREENHOUSE GAS ENERGY

- Relatively clean source of electricity
 - Suitable for large power generation
- But what about the hydrogen economy?
 - Fuel cells for cars (gallon gas = 5 lbs C)
- A primary energy source is needed to produce H_2
 - Hydrogen is a *energy carrier*
- Need to solve the entire cycle




FUSION RESEARCH IS INTERNATIONAL IN SCOPE



- Open international research


- Long history of collaboration



- Significant progress

- Ready for the next step

ITER IS THE NEXT STEP FOR FUSION



- ~\$5B class device, over 20 countries
 - Number 1 DOE/SC Facility Priority
- Built in either France or Japan
 - International collaboration
- Successful operation requires
 - Large simulations, shared vis, decisions back to the control room
- How does the U.S. benefit?
 - Effective remote collaboration!

THE NATIONAL FUSION COLLABORATORY PROJECT

- Funded DOE Office of Advanced Scientific Computing Research
 - Part of the SciDAC initiative
- Unify distributed MFE research into a U.S. Virtual Organization

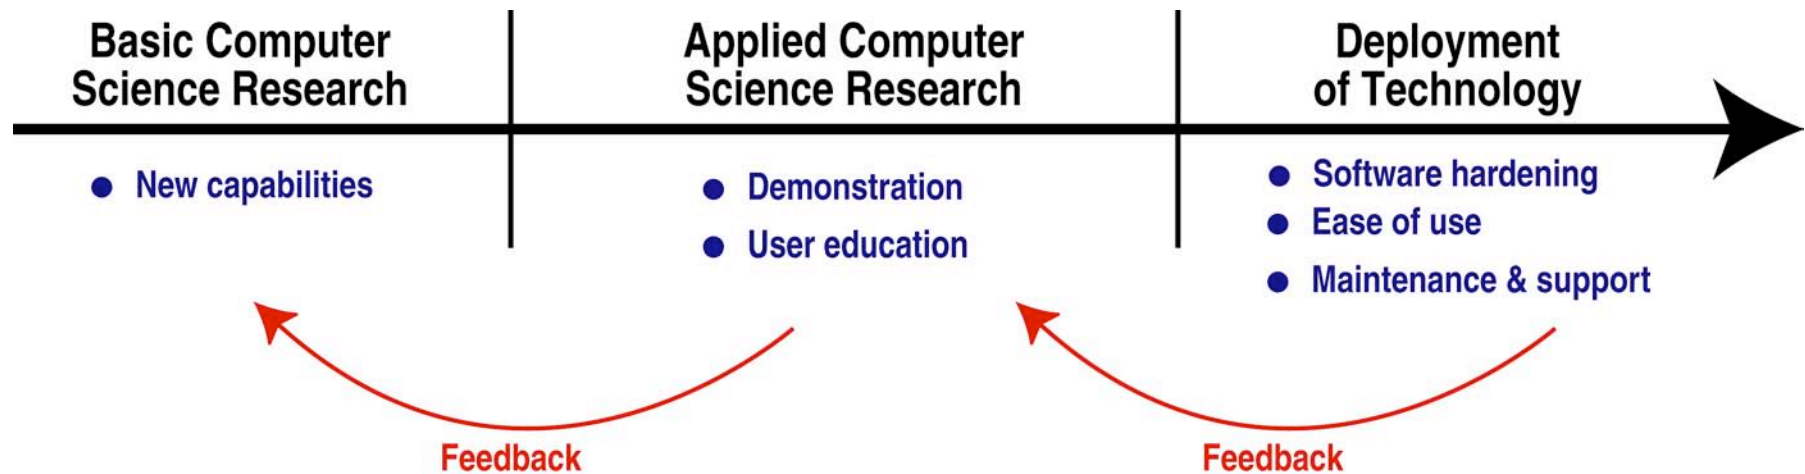
NFC'S GOALS

- More efficient use of experimental facilities
- Integrate theory & experiment
- Facilitate multi-institution collaboration
- Create standard tool set




THE NFC PROJECT BENEFITS FROM A DIVERSE TEAM

Synergistic benefits derived from interdisciplinary interactions
& indirect interactions with other SciDAC projects



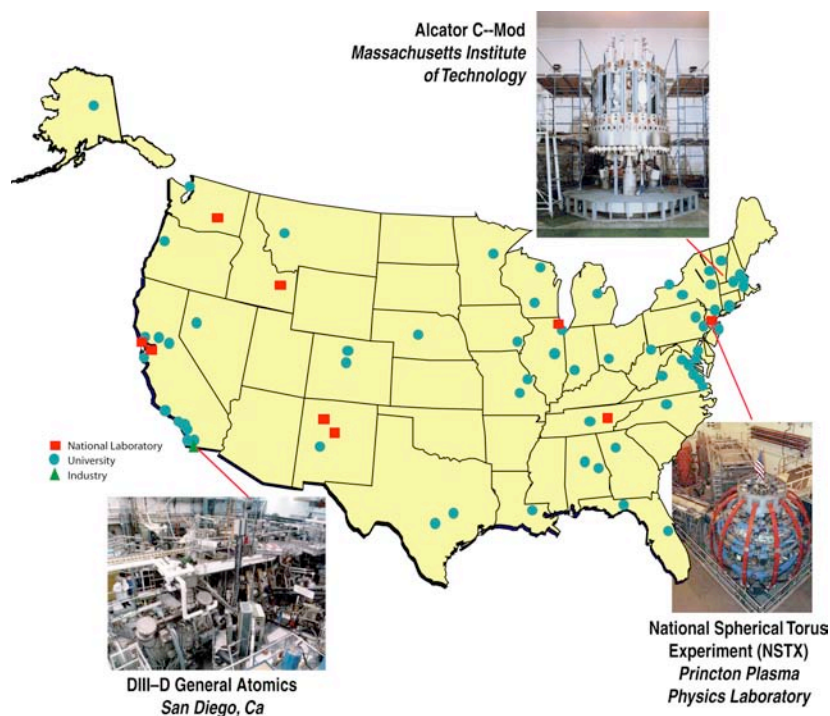
- ANL: Distributed Systems Lab
- ANL: Futures Lab
- General Atomics: DIII-D Fusion Lab
- LBNL: Distributed Systems

- MIT: C-Mod Fusion Lab
- Princeton Computer Science
- PPPL: NSTX Fusion Lab
- Utah: Scientific Computing & Imaging



THREE LARGE U.S. EXPERIMENTAL FACILITIES AND A VIBRANT THEORETICAL COMMUNITY

Collaboratory is required to advance fusion science: geographically diverse community (37 states, 3 large experiments), leading to 1 worldwide experiment



- 3 Large Experimental Facilities
 - ~\$1B replacement cost
- 40 U.S. fusion research sites
 - Over 1500 scientists
- Efficient collaboration is required!
 - Integrate geographically diverse teams
- One future worldwide machine
 - Not based in US
 - US needs collaboration tools to benefit

EXPERIMENTAL SCIENCES PLACES A LARGE PREMIUM ON RAPID DATA ANALYSIS IN NEAR-REAL-TIME



DIII-D Control Room

- Pulsed experiments
 - 10s duration plasma every 20 minutes
- 20-40 people in control room
 - More from remote locations
- 10,000 separate measurements/plasma
 - kHz to MHZ sample rates
 - Between pulse analysis
- Not batch analysis and not a needle in a haystack problem
 - Rapid “real-time” analysis of many measurements
- More informed decisions result in better experiments
 - The collaborative control room


THE VISION FOR THE NFC'S TECHNOLOGIES

- Data, Codes, Analysis Routines, Visualization Tools should be thought of as network accessible services
 - Access is stressed rather than portability
 - Transparency and ease of use are crucial elements
 - Not CPU cycle scavenging or “distributed” supercomputing
- Shared security infrastructure with distributed authorization and resource management
 - Ease of use: “security with transparency”
 - X.509 certificates from a trusted Certificate Authority
 - Distributed authorization allows stakeholders to control their resources
- Collaborative nature of research requires shared visualization applications and widely deployed collaboration technologies
 - Integrate geographically diverse groups

Optimize the most expensive resource - people's time


NFC'S WORK FOCUSES ON FOUR AREAS

- Production Grid computational Service
 - Not CPU cycle scavenging & dynamic configuration
 - Move high-performance distributed computing to the WAN
- Compare theory & experiment: large datasets
 - Both a visualization and data handling challenge
- Fully engaged remote scientists
 - Real-time human interaction at a distance
- Tiled display wall for collocated collaboration
 - For the large groups in the tokamak control room

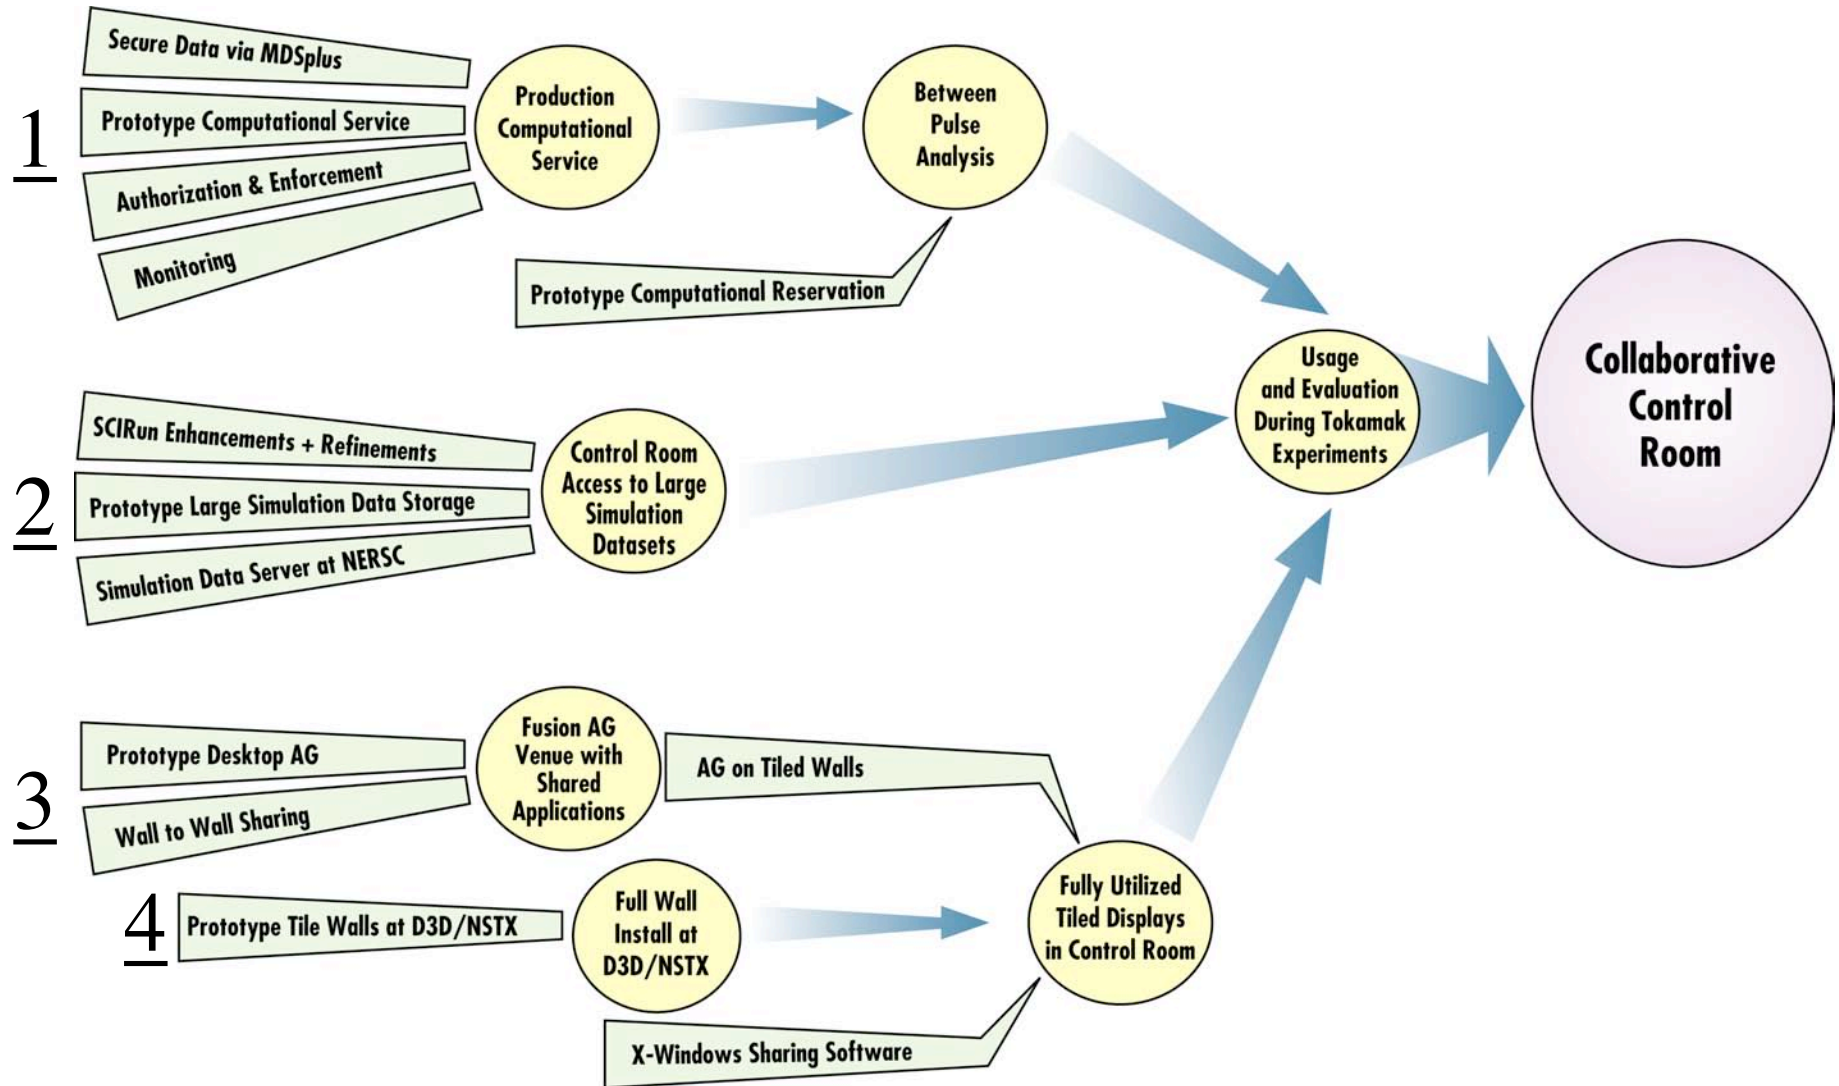


THE COLLABORATIVE CONTROL ROOM IS FUNDAMENTAL TO ADVANCING FUSION SCIENCE

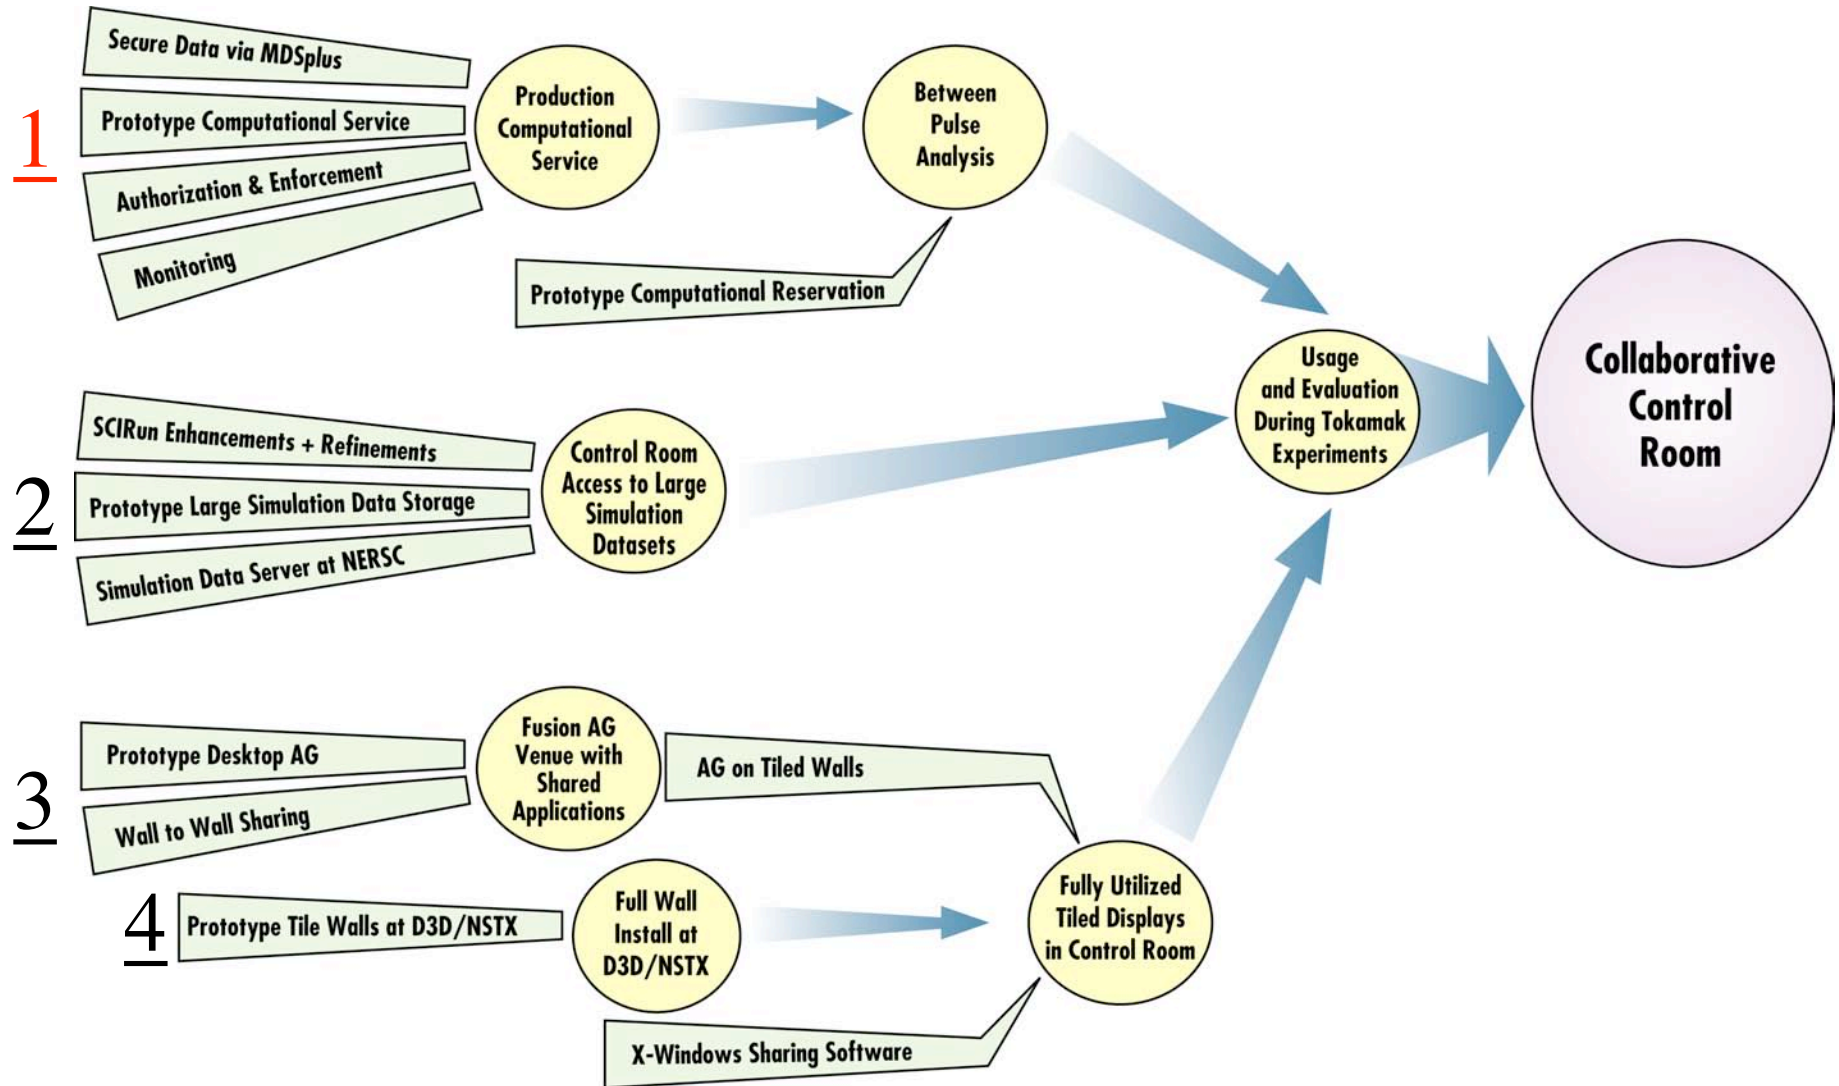
- Secure computational resources that can be scheduled as required
- Rapidly compare experimental data to simulation results
- Share individual results with the group via shared displays
- Fully engaged remote scientists with audio, video, shared displays



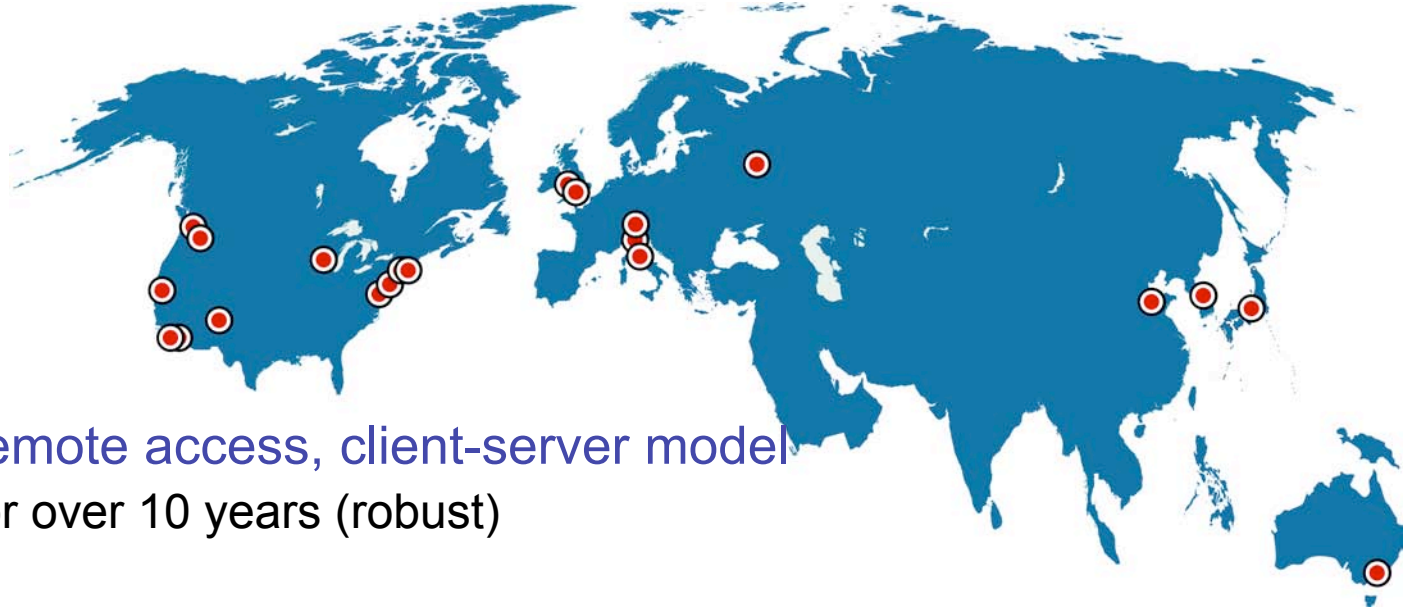
WORK TOWARDS THE COLLABORATIVE CONTROL ROOM



WORK TOWARDS THE COLLABORATIVE CONTROL ROOM




FUSIONGRID: SECURE ACCESS TO FUSION DATA



- MDSplus: remote access, client-server model
 - In use for over 10 years (robust)
- Wide adoption worldwide
 - Unifying the data interface (e.g. Visualization)
- MDSplus data access secured with Globus GSI
 - X.509 certificates
- MyProxy used to manage certificates
 - FusionGrid CA created

ROAM – RESOURCE ORIENTED AUTHORIZATION MANAGER FOR FUSIONGRID

- Centralized authorization of Grid resources
 - Resource provider can implement simple or complex policies
- Anyone can create resource and grant users permissions
 - The site is always a resource (for site-security officer)
- User identity based on FusionGrid certificate DN
 - Supports “linked” DN’s for users with multiple certificates
- Automatic ROAM account creation on first access
 - Secure access with certificate or MyProxy login
- Handles dynamic accounts!
 - Globus Gridmap callouts can replace grid-mapfile
- XIO http driver used to check resource permissions



ROAM HAS AN EASY TO USE WEB INTERFACE

User's view of their resource permissions

Resource	Permission	Description	Action	Logs
CMOD Data	Write	Access to CMOD experiment data	Submit Inquiry	Show log
CMOD Data	Admin	Access to CMOD experiment data	Administer Resource	Show log
CMOD-Jobs	Admin	Permission to execute CMOD-Jobs	Administer Resource	Show log
DIID-MDSplus	Read	MDSplus at GA	Submit Inquiry	Show log
DIID-MDSplus	Admin	MDSplus at GA	Administer Resource	Show log

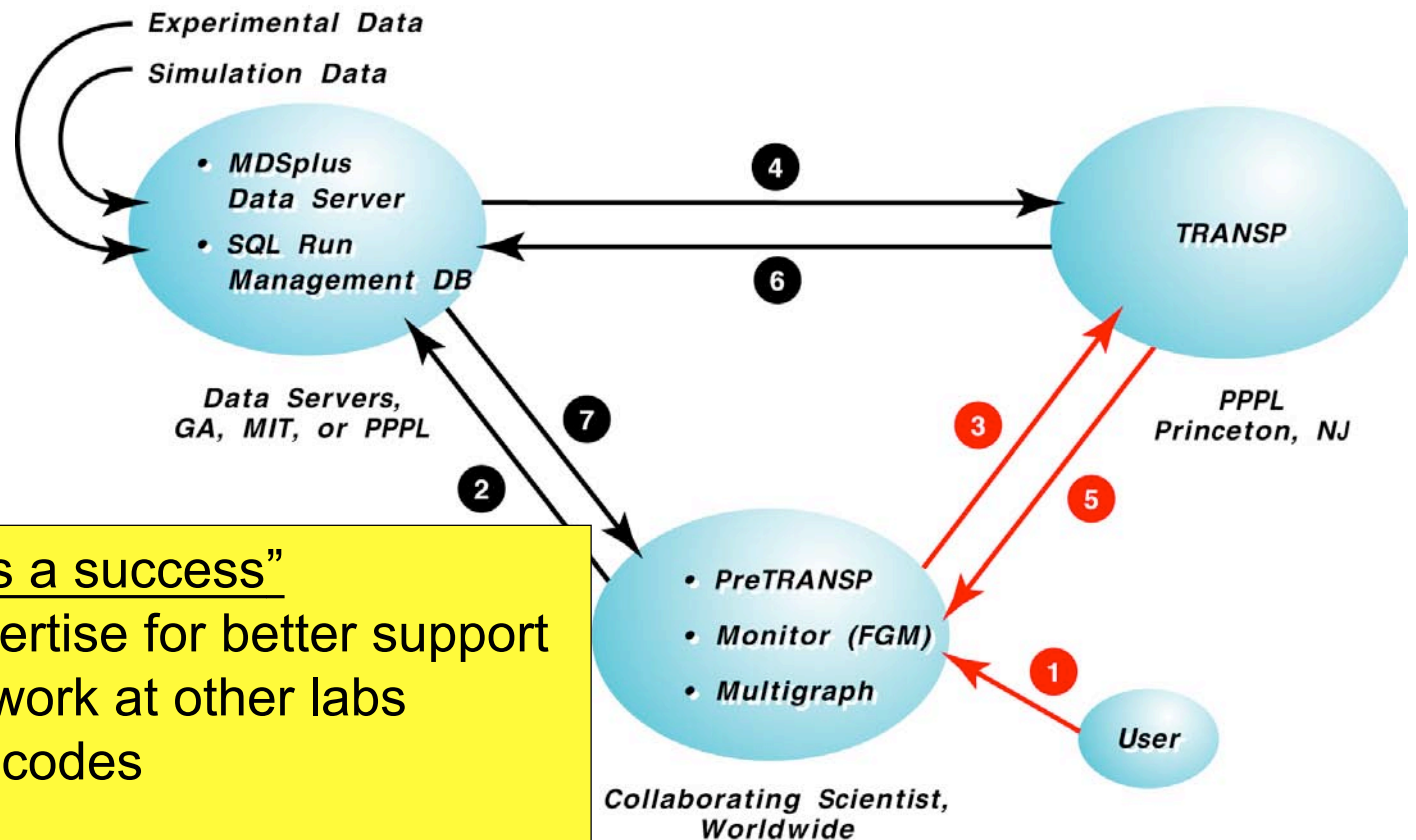
Resource's view of what users have permission

Resource Authorizations					Add New
User	Permission	Auxilliary			
Burruss, Justin	Execute	burruss	Modify	Delete	
Burruss, Justin	Admin		Modify		
Clamage, David	Execute	clamage	Modify	Delete	
Flanagan, Sean	Execute		Modify	Delete	
Fredian, Thomas	Execute	fredian	Modify	Delete	
Greenwald, Martin	Execute		Modify	Delete	
Qian, Min	Execute		Modify	Delete	
Stillerman, Joshua	Execute		Modify	Delete	
Thomas, Mary	Execute		Modify	Delete	

Example New Resource

- Create new resource A
- Create new resource Site-A
- Create policy for A
- Grant Admin privilege for Site-A to site administrator
- Grant Access permission to user for Site-A
- Grant Execute permission to user for A

SUCCESSFUL GRID COMPUTING FOR FUSION SCIENCE



"This is a success"

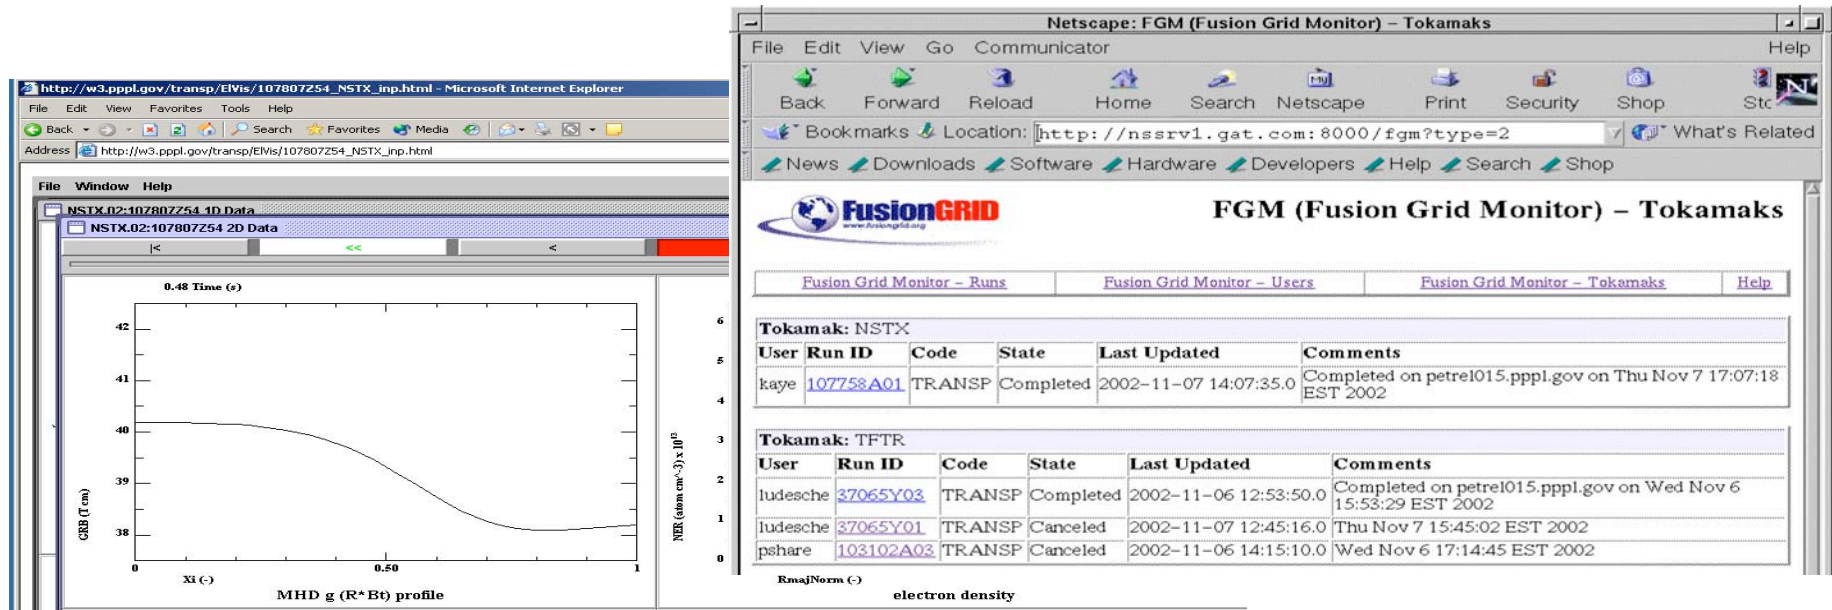
- Centralized expertise for better support
- Reduce admin work at other labs
- Model for other codes

The U.S. TRANSP Service

- 4,200 cases, 10,000 CPU hrs
- 10 fusion experimental machines



FUSION GRID MONITOR: AN EFFICIENT APPLICATION MONITORING SYSTEM FOR THE GRID ENVIRONMENT

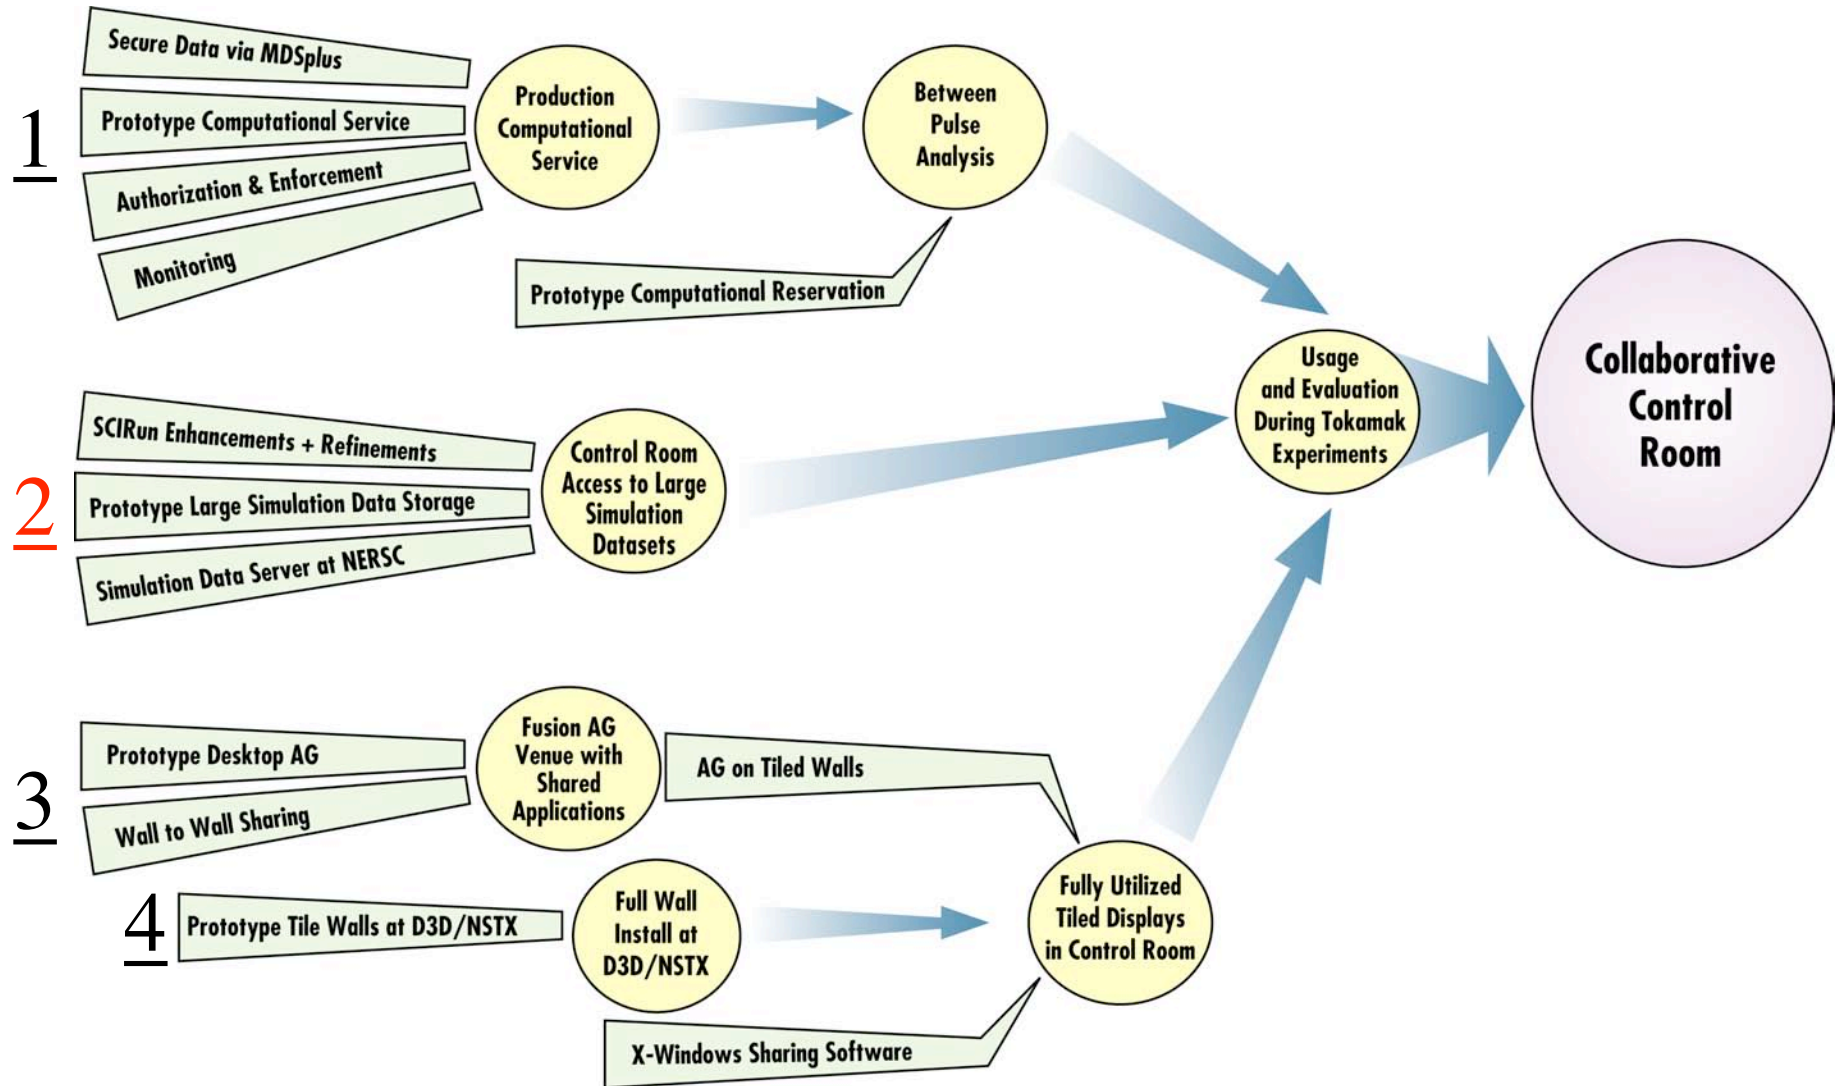


- Users track and monitor the state of applications on FusionGrid
 - Output dynamically via HTML, Built as Java Servlet (JDK2.1)
- Code maintenance notification
 - Users notified, queuing turned off, code rebuilt, queue restarted
- Results of simulation visualized during run
 - Both input and output quantities

ADVANCED RESERVATION COMPUTATION FOR DATA ANALYSIS TO SUPPORT EXPERIMENTAL SCIENCE


- Long-term vision: SciDAC code on supercomputer between pulses
 - Data management
 - Network QoS
 - Visualization
 - CPU scheduling
 - Faster CPUs and algorithms
- End-to-end agreement being prototyped in the NFC project
 - CPU reservation
 - Network transfer agreements based on simple prediction
- FusionGrid service TRANSP will be tested between pulses
 - First such capability for FES research

WORK TOWARDS THE COLLABORATIVE CONTROL ROOM



SCIRUN TO VISUALIZE COMPLEX SIMULATIONS FOR BETTER UNDERSTANDING

- Open source, multi-platform capable for a wide user base
- To facilitate quantitative comparison of simulations & experimental results

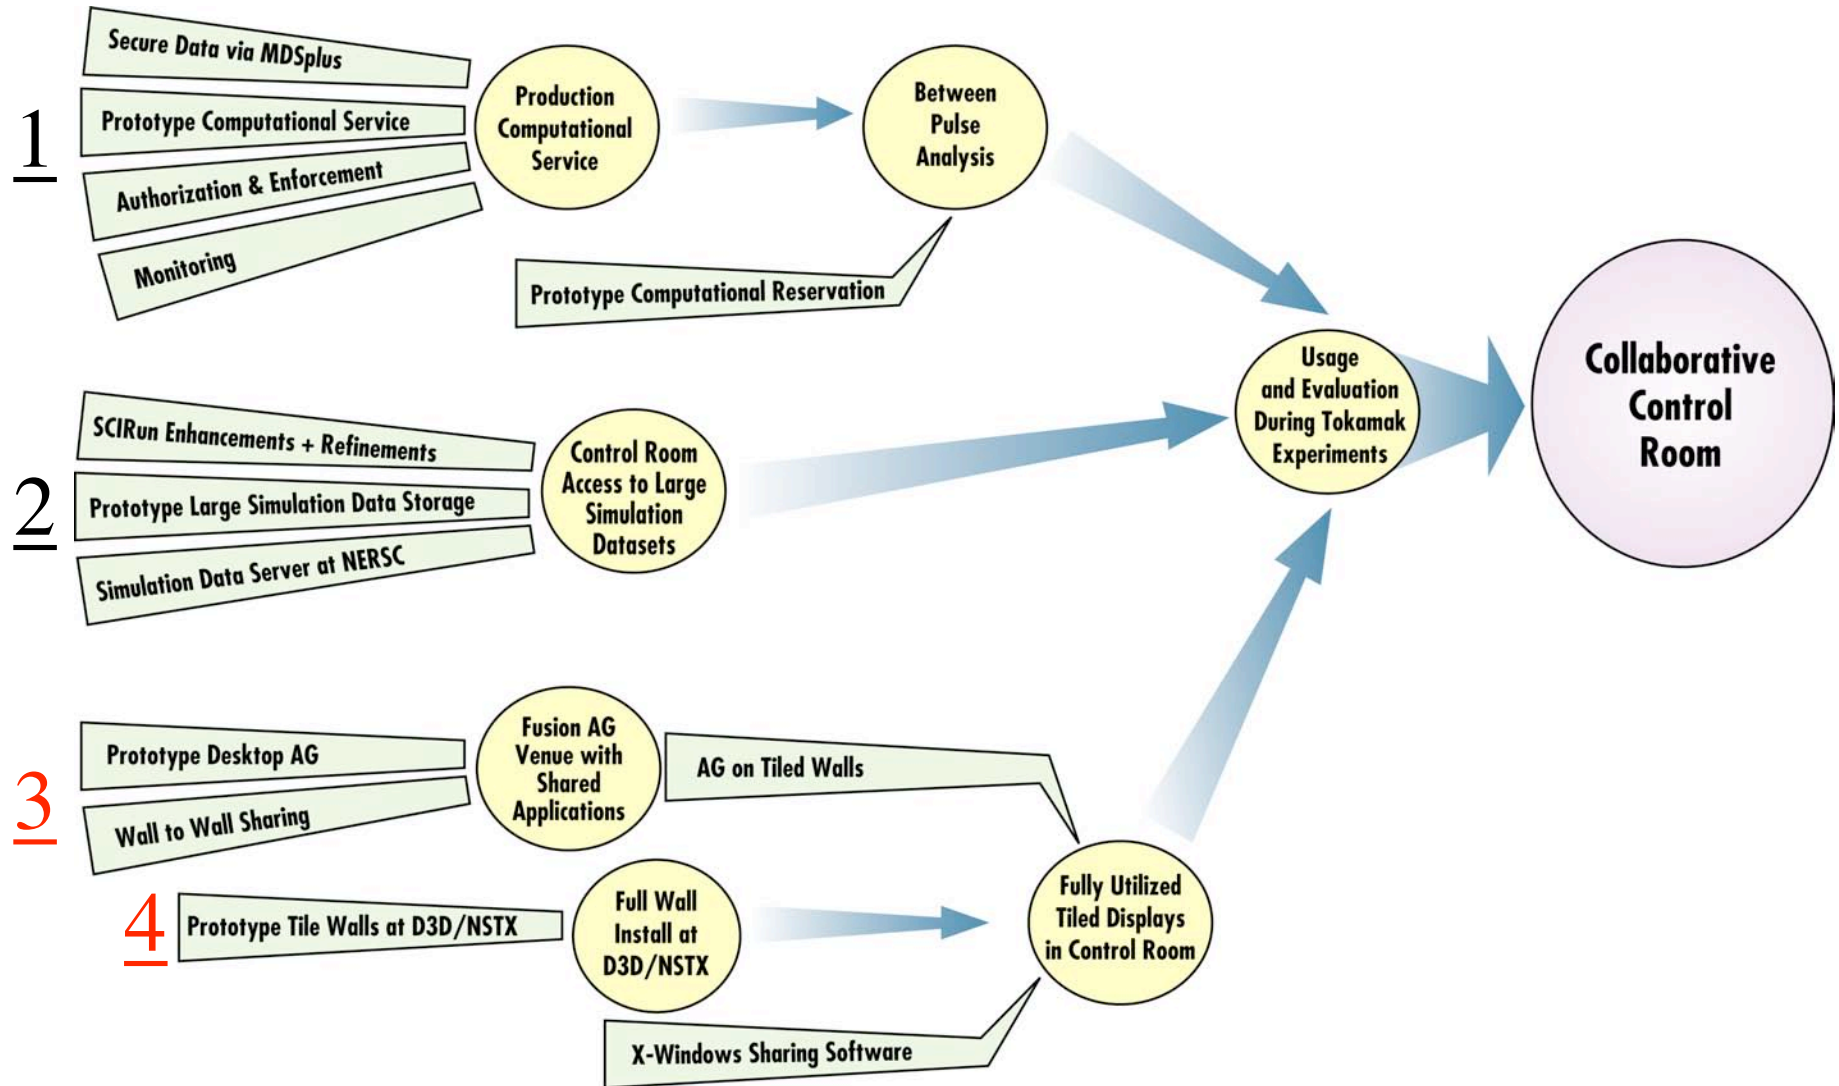


SciDAC CEMM NIMROD Simulation of a DIII-D Plasma

Raising the challenge
of very large datasets

- MDSplus
- Storage method
- Data location
- Parallel I/O

WORK TOWARDS THE COLLABORATIVE CONTROL ROOM



TILED DISPLAYS INSTALLED IN FUSION CONTROL ROOMS

DIII-D Tokamak Control Room



NSTX Tokamak Control Room



- Enhanced collaboration within a large control room
 - “Publish” your analysis for the group to see and discuss
- Share and collaborate between tiled displays
 - Clone of tokamak control room (discussed for ITER)

ACCESS GRID: REAL TIME COMPLEX COMMUNICATION



- Small to large immersive nodes
 - Open source software
 - Commodity hardware
 - Software framework for Collab.
 - Shared applications
 - Customized software allowed
-
- Being used for seminars, working meetings, tokamak operations
 - AG nodes are international: 300 nodes and counting
 - Mac OS X version available soon


Interoperability of technologies is critical

SC03 DEMO: COLLABORATIVE CONTROL ROOM

Shared Application

Video & Audio

Real Time Data Display



SuperComputing 2003, Phoenix AZ

THE NFC IS BEING SUCCESSFUL FOR FUSION SCIENCE

- Success is not defined by demos or papers published
 - Although these are important aspects of the project
- Success is defined by positive impact on the science
 - True metric is the actions of the fusion scientists
- A main fusion code is only available via NFC Grid computing
- The fusion program is spending its own money to purchase equipment and to support the usage of new NFC functionality
- A fusion scientist decided to remain in San Diego and use NFC technology to lead the fusion experiment in England

There is no going back:
SciDAC and the NFC Project has changed
FES collaborative culture

USING NFC TECHNOLOGIES TO UNDERSTAND THE COLLABORATIVE CONTROL ROOM

January 2004, San Diego



Scientific Leadership of the JET
Tokamak in England from the US

DIII-D Tokamak Control Room - July 2003




Remote Participation into the
DIII-D Control Room


- Also for Japan - US & US-Germany collaborations
- Technology applicable to smaller collaborations; diagnostic control

NFC TECHNOLOGIES SCALE TO THE NEXT DEVICE

KSTAR



ITER



- One physical location
 - International collaboration
- Pulsed experiment with simulations
 - ~TBs of data in 30 minutes
- Successful operation requires
 - Large simulations, shared vis, decisions back to the control room
 - Remote Collaboration via FusionGrid
- NFC technology being considered as the model for ITER & KSTAR

LESSONS LEARNED AND OUTSTANDING ISSUES

- Difficulties combining Grid-security and Site-security (firewalls)
 - Greatly limiting the potential expansion of the user base
- Certificate management for users and developers too difficult
 - This is often their first experience: needs to be positive
- Manipulating large multi-dimensional datasets is still a challenge
 - Need to test new approaches
- Control room presence is more than audio/video & shared apps
 - Include things one sees & hears when physically in control room
- Difficulty combining H.323, VRVS, and AG into one united meeting
 - Unlikely one solution is adopted, need easy interplay between systems
- Education, training and shared documentation
 - Users, developers, and system administrators
- Continue to learn and interact with other scientific domains
 - No need to reinvent the wheel

CONCLUDING COMMENTS

- The National Fusion Collaboratory Project is implementing and testing new collaborative technologies for fusion research
 - FusionGrid services being used to benefit daily FES research
- Clear vision forward to the collaborative control room
 - Concept encompasses most if not all collaborative FES needs
- Collaborative technology critical to the success of the FES program
 - Experimental: Fewer, larger machines in future (ITER)
 - Computation: Moving toward integrated simulation (FSP)

First on our list is fusion. The prospect of limitless source of clean energy for the world leads with our commitment to join the international fusion energy experiment known as ITER.

— Secretary of Energy Spencer Abraham, November 10, 2003

Introducing the Department's 20-year plan for building the scientific facilities of the future.